UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/659,374	09/11/2003	Kentaro Shimada	H-1112	3004

24956 7590 01/30/2006

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.
1800 DIAGONAL ROAD
SUITE 370
ALEXANDRIA, VA 22314

EXAMINER	
DARE, RYAN A	

ART UNIT	PAPER NUMBER
2186	

DATE MAILED: 01/30/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/659,374	SHIMADA, KENTARO	
	Examiner	Art Unit	
	Ryan Dare	2186	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *14 December 2005*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-15* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-15* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *14 December 2005* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☒ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☒ All b)☐ Some * c)☐ None of:

1.☒ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: _____.

DETAILED ACTION

Information Disclosure Statement

1. Examiner was able to locate Applicant's copy of the IBM TotalStorage Enterprise Storage Server Model 800 reference, and was able to confirm that this is in fact, the same reference that Examiner has cited and considered.

Drawings

1. The replacement drawing sheets filed on 12/14/05 are approved.

Specification

2. The amendments to the specification filed on 12/14/05 are approved. The objections to the minor informalities in the specification are withdrawn.

Claim Objections

2. The amendments to claim 1 filed on 12/14/05 is approved. The objection to the typographical error has been withdrawn.

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the

invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

4. The factual inquiries set forth in *Graham* **v.** *John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

5. Claims 1, 10 and 15 are rejected under 35 U.S.C. 103(a) as being unpatentable over the IBM publication *IBM TotalStorage Enterprise Storage Server Model 800*, hereafter IBM, in view of US Patent 5,123,099 issued to Shibata et al, further in view or Kinjo et al., US Patent 6,944,684.

6. With respect to claim 1, IBM discloses:

A storage system comprising:

a storage device, in page 101, figure 4-7;

and a first controller and a second controller both connected to the storage device and a host system, in page 101, figure 4-7, where "Cluster 1" is the first controller and "Cluster 2" is the second controller and the host is connected to one of the host adapters labeled "HA".

IBM discloses a system where the data from the host system is written into a cache of a first controller and written to nonvolatile memory of the second controller. The host is notified that the I/O operation is complete as soon as the data is written.

See page 89 of the IBM publication and page 12 of Applicant's Petition to Make Special. However, IBM fails to teach the system where data is written to two memories on the first controller, before transferring the data to the memory on the second controller.

Shibata et al. describe another backup storage system similar in scope to the IBM reference. Shibata et al. also describe a storage system with two controllers,

wherein, the first controller has a first memory, said first memory being a cache memory in fig. 1, numeral 3, and a second memory, in fig. 1, numeral 4,

wherein, the second controller has a third memory, in fig. 1, numeral 2',

wherein, in the case where the first controller receives data from the host system, the first controller stores the data in the first and the second memories, and then the first controller transfers the data stored in the second memory to the third memory, in col. 4, lines 31-38.

Shibata et al. writes the data into two memories on the first controller, so it is possible to transfer the contents of the second memory to the third memory without an operation of the CPU of the first controller. Shibata et al. describe the performance advantages of this configuration in column 4, lines 39-66. Among them is that the first and second memories can be written to at the same speed. With the system of IBM, the CPU of controller 1 may have to worry about delays due to cross connections at the time of data writing to both of the main storage devices. The system of Shibata et al. eliminates this concern and uses the high speed of the cache memories properly.

Shibata fails to teach that the third memory is a cache memory that receives a copy of the data received by said first memory. The third memory of Shibata et al. is a

main memory unit. Kinjo et al. teaches that the memory on the second controller can be a cache memory (i.e. the third memory), in col. 6, lines 6-11. This teaches the limitation:

wherein, the second controller has a third memory, said third memory being a cache memory that receives a copy of the data received by said first memory.

7. It would be obvious to one of ordinary skill, having the teachings of IBM and Shibata et al. before him at the time the invention was made, to combine the backup storage system of IBM with the backup storage system of Shibata et al. to eliminate the performance problems due to cross connectional delays, as taught by Shibata et al., and to include response messages as taught by IBM et al. so the host knows when the data has been written to the duplex memories.

Furthermore, it would be obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata, and Kinjo et al. before him at the time the invention was made, to modify the backup storage system of IBM and Shibata with the backup storage system of Kinjo et al. in order to allow high-speed data copy between controllers, as taught by Kinjo et al. in col. 8, lines 61-65. The prior art system of Shibata possibly suffers performance penalties due to the fact that the third memory is a main memory, so the system of Kinjo et al. resolves this deficiency by using a high-speed cache memory as a third memory.

8. With respect to claim 10, Shibata et al. disclose:

A disk drive, in page 101, figure 4-7;

A first cache memory for storing data sent from a host system so as to be written on the disk drive, in the cache of Cluster 1, shown on page 101, figure 4-7;

A second cache memory for storing a duplicate of the data to be written on the disk drive, in the NVS of Cluster 2, shown on page 101, figure 4-7;

In addition, the IBM reference also teaches informing the host system that the data writing is complete. See page 89 of the IBM reference and page 12 of Applicant's Petition to Make Special. However, IBM fails to teach the system where data is written a first memory and a FIFO buffer on the first controller, before transferring the data to the memory on the second controller.

Shibata et al. describe another backup storage system similar in scope to the IBM reference. Shibata et al. describe a storage system with a FIFO buffer for temporarily storing the duplicate of the data sent from the host system to transfer the duplicate of the data to a second memory, in col. 4, lines 39-44.

Shibata et al. writes the data into a first memory and a FIFO buffer, so it is possible to transfer the contents of the FIFO buffer to the second memory without an operation of the CPU of the first controller. Shibata et al. describe the performance advantages of this configuration in column 4, lines 39-66. Among them is that the first memory and FIFO buffer can be written to at the same speed. With the system of IBM, the CPU of controller 1 may have to worry about delays due to cross connections at the time of data writing to both of the main storage devices. The system of Shibata et al. eliminates this concern and uses the high speed of the cache memories properly.

Shibata fails to teach that the memory on the second controller is a cache memory that receives a copy of the data received by said first memory. The memory on the second controller of Shibata et al. is a main memory unit. Kinjo et al. teaches that the memory on the second controller can be a cache memory in col. 6, lines 6-11. This teaches the limitation:

wherein, the second controller has a third memory, said third memory being a cache memory that receives a copy of the data received by said first memory.

9. It would be obvious to one of ordinary skill in the art, having the teachings of IBM and Shibata et al. before him at the time the invention was made, to combine backup storage system of IBM with the backup storage system of Shibata et al. to eliminate the performance problems due to cross connectional delays, as taught by Shibata et al., and to include response messages as taught by IBM et al. so the host knows when the data has been written to the duplex memories.

Furthermore, it would be obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata, and Kinjo et al. before him at the time the invention was made, to modify the backup storage system of IBM and Shibata with the backup storage system of Kinjo et al. in order to allow high-speed data copy between controllers, as taught by Kinjo et al. in col. 8, lines 61-65. The prior art system of Shibata possibly suffers performance penalties due to the fact that the third memory is a main memory, so the system of Kinjo et al. resolves this deficiency by using a high-speed cache memory as the memory on the second controller.

10. With respect to claim 15, Shibata et al. disclose:

A method of writing data in a storage system having duplex cache memory, comprising the steps of:

writing the data in one of the duplex cache memory for duplicating and storing data sent from the host system, in fig. 1 and par. 4, lines 31-38;

writing the data in a FIFO buffer capable of performing writing with a speed higher than that of the other one of the duplex cache memory, in par. 4, lines 31-44. Note that in the system of Shibata et al., the FIFO buffer is written to at the same speed as the cache memory on the first controller. In the system of IBM, the duplex cache is on the second controller. This can introduce delays due to cross connections as discussed by Shibata et al. in col. 4, lines 44-55. The system of Shibata et al. is able to write this data faster to the FIFO than it would be if the duplex cache were on the second controller as with the IBM reference.

Shibata et al. fail to teach a response to the host system. IBM teaches confirming that the data have been correctly written in the duplex cache memories and informing the host system about completions of the data writing. See page 89 of the IBM reference and Applicant's Petition to Make Special.

Shibata et al. teach writing the data written in the FIFO buffer to a memory on the second controller, in col. 4, lines 31-38.

Shibata fails to teach that the memory on the second controller is a cache memory that receives a copy of the data received by said first memory. The memory on the second controller of Shibata et al. is a main memory unit. Kinjo et al. teaches that

the memory on the second controller can be a cache memory in col. 6, lines 6-11. This teaches the limitation:

wherein, the second controller has a third memory, said third memory being a cache memory that receives a copy of the data received by said first memory.

11. It would be obvious to one of ordinary skill in the art, having the teachings of IBM and Shibata et al. before him at the time the invention was made, to combine backup storage system of IBM with the backup storage system of Shibata et al. to eliminate the performance problems due to cross connectional delays, as taught by Shibata et al., and to include response messages as taught by IBM et al. so the host knows when the data has been written to the duplex memories.

Furthermore, it would be obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata, and Kinjo et al. before him at the time the invention was made, to modify the backup storage system of IBM and Shibata with the backup storage system of Kinjo et al. in order to allow high-speed data copy between controllers, as taught by Kinjo et al. in col. 8, lines 61-65. The prior art system of Shibata possibly suffers performance penalties due to the fact that the third memory is a main memory, so the system of Kinjo et al. resolves this deficiency by using a high-speed cache memory as the memory on the second controller.

12. Claims 2-7 and 11-13 are rejected under 35 U.S.C. 103(a) as being unpatentable over IBM, Shibata et al. and Kinjo et al. as applied to claims 1, 10 and 15 above, and further in view of Beardsley et al., US Patent 5,437,022.

13. With respect to claim 2, IBM, Shibata et al. and Kinjo et al. teach all other limitations of the parent claim as discussed supra, but fail to teach the power system of claim 2. Beardsley et al. disclose:

A storage system according to claim 1, further comprising a first power unit and a second power unit, numerals 80 and 90 of fig. 4,

wherein the first memory of the first controller receives power feeding from the first power unit, in fig. 4, with cache A being the first memory and numeral 80 being the first power unit,

wherein the second memory of the first controller and the third memory of the second controller receive power feeding from the second power unit, in fig. 4, with NVS A being the second memory of the first controller, and Cache B being the third memory of the second controller, and numeral 90 being the second power unit.

14. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al., and Beardsley et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al. and Kinjo et al. to include a system of power management as taught by Beardsley et al. With respect to the IBM reference, the first cache memory and the backup nonvolatile storage are on separate controllers with separate power supplies. In case one power supply fails, you have a backup on the other system, as taught on page 30, section 2.7.3 of the IBM reference. The backup storage system of Shibata et al. more closely resembles the duplex cache writing of the present invention, but fails to mention power management. It would have been obvious to use the system of Beardsley et al. so if a

failure in power occurs at the first power source, you will still have a backup of the data in the other memories, and won't lose data, as taught by Beardsley et al. in col. 6, lines 3-8.

15. With respect to claim 3, IBM, Shibata et al., and Kinjo et al. teach all other limitations of the parent claim as discussed supra, but fail to teach the power system of claim 3. Beardsley et al. disclose:

A storage system according to claim 1, further comprising a first, a second and a third power unit, in fig. 4, numerals 82, 96 and 92 being the first, second and third power units respectively,

wherein the first memory receives power feeding from the first power unit, in fig. 4, numerals 58 and 82,

wherein the second memory receives power feeding from the second power unit, in fig. 4, numerals 60 and 96.

wherein the third memory receives power feeding from the third power unit, in fig. 4, numerals 61 and 92.

16. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al., and Beardsley et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al., and Kinjo et al. to include a system of power management as taught by Beardsley et al. With respect to the IBM reference, the first cache memory and the backup nonvolatile storage are on separate controllers with separate power supplies. In case one power supply fails, you have a backup on the other system, as taught on page 30, section

2.7.3 of the IBM reference. The backup storage system of Shibata et al. more closely resembles the duplex cache writing of the present invention, but fails to mention power management. It would have been obvious to use the system of Beardsley et al. so if a failure in power occurs at any one of the power sources, you will still have a backup of the data in the other memories, and won't lose data, as taught by Beardsley et al. in col. 6, lines 3-8.

17. With respect to claim 4, IBM, Shibata et al., and Kinjo et al. teach all other limitations of the parent claim as discussed supra, but fail to teach the power system of claim 4. Beardsley et al. disclose:

A storage system according to claim 1, further comprising a power unit,

wherein the first memory and the second memory receive power feeding from the power unit, in fig. 4, with numeral 58 and numeral 59 being the first and second memories, and numeral 80 being the first power unit.

Beardsley et al. fail to teach a battery for the second memory. IBM teaches a battery for the second memory on page 89. IBM also teaches a battery charging system on the NVS card, where the battery is charged by the power unit, on page 30, paragraph 4.

18. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al., and Beardsley et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al. and Kinjo et al. to include a system of power management as taught by Beardsley et al. With respect to the IBM reference, the first cache memory and the backup nonvolatile

storage are on separate controllers with separate power supplies. In case one power supply fails, you have a backup on the other system, as taught on page 30, section 2.7.3 of the IBM reference. The backup storage system of Shibata et al. more closely resembles the duplex cache writing of the present invention, but fails to mention power management. It would have been obvious to use the system of Beardsley et al. so if a failure in power occurs at the first power source, you will still have a backup of the data in the other memories, and will not lose data, as taught by Beardsley et al. in col. 6, lines 3-8. Furthermore, it would have been obvious to one of ordinary skill to also include a battery backup for the FIFO buffer, so you do not lose data in the event of a power failure. In the storage system of IBM, it would be obvious to use the power unit to charge the battery, to prevent data loss, as taught by IBM on page 51, paragraph 2.

19. With respect to claim 5, Shibata et al., IBM, Kinjo et al. and Beardsley et al. teach all other limitations of the parent claims as discussed supra, and IBM teaches:

A storage system according to claim 4,

wherein, in the case where a failure occurs in the power unit, the second memory switches the power feeding from the power unit to power feeding from the battery, on page 89, paragraph 4 where it says "The battery will power the NVS for up to 72 hours following a total power failure."

20. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al. and Beardsley et al. before him at the time the invention was made to modify the backup storage system of Shibata et al., Kinjo et

al. and Beardsley et al. to include a battery to switch to in case of power failure in order to prevent data loss as taught by IBM on page 51, paragraph 2.

21. With respect to claim 6, IBM, Shibata et al., Kinjo et al., Beardsley et al. teach all other limitations of the parent claims as discussed supra, but IBM, Beardsley et al. and Kinjo et al. fail to teach that the second memory is a FIFO buffer. Shibata et al. disclose that the second memory can be a FIFO buffer in col. 4, lines 39-44 and in fig. 1.

22. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al., and Beardsley et al. before him at the time the invention was made to modify the backup storage system so that the second memory is a FIFO buffer to eliminate the performance problems due to cross connectional delays, as taught by Shibata et al. in col. 4, lines 39-55, since the FIFO can copy the data over to the other controller independently of the processing unit.

23. With respect to claim 7, IBM, Shibata et al., Kinjo et al. and Beardsley et al. teach all other limitations of the parent claims as discussed supra, but Shibata et al. fail to teach that the storage device can be a plurality of storage devices.

IBM shows that the storage device can be a plurality of storage devices in a RAID configuration as shown in fig. 1 on page 50.

24. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al. and Beardsley et al. before him at the time the invention was made to modify the backup storage system so that the storage device is a plurality of storage devices which may have redundant storage, which prevents data loss, as taught by IBM on page 53, section 3.3.1.

25. With respect to claim 11, IBM, Shibata et al., and Kinjo et al. teach all other limitations of the parent claim as discussed supra, but fail to teach the power system of claim 11. Beardsley et al. disclose:

A storage system according to claim 10, further comprising:

a first power unit connected to the first cache memory, in fig. 4, numerals 58 and 82;

and a second power unit connected to the second cache memory, the second power unit being independent from the first power unit, in fig. 4, numerals 60 and 96,

Beardsley teaches a third cache memory being connected to the second power unit in fig. 4, numerals 61 and 92, but fails to teach a FIFO buffer. Shibata et al. teaches that the third memory can be a FIFO buffer, in col. 4, lines 39-44.

26. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al., and Beardsley et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al., and Kinjo et al. to include a system of power management as taught by Beardsley et al. With respect to the IBM reference, the first cache memory and the backup nonvolatile storage are on separate controllers with separate power supplies. In case one power supply fails, you have a backup on the other system, as taught on page 30, section 2.7.3 of the IBM reference. The backup storage system of Shibata et al. more closely resembles the duplex cache writing of the present invention, but fails to mention power management. It would have been obvious to use the system of Beardsley et al. so if a failure in power occurs at any one of the power sources, you will still have a backup of

the data in the other memories, and won't lose data, as taught by Beardsley et al. in col. 6, lines 3-8.

27. With respect to claim 12, IBM, Shibata et al. and Kinjo et al. teach all other limitations of the parent claim as discussed supra, but fail to teach the power system of claim 11. Beardsley et al. disclose:

A storage system according to claim 10, further comprising:

a first power unit connected to the first cache memory, in fig. 4, numerals 58 and 82;

a second power unit connected to the second cache memory, the second power unit being independent from the first power unit, in fig. 4, numerals 60 and 96,

Beardsley et al. teach a third cache memory being connected to the second power unit, the third power unit being independent from the first power unit, in fig. 4, numerals 61 and 92, but fails to teach a FIFO buffer. Shibata et al. teaches that the third memory can be a FIFO buffer, in col. 4, lines 39-44.

28. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al. and Beardsley et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al. and Kinjo et al. to include a system of power management as taught by Beardsley et al. With respect to the IBM reference, the first cache memory and the backup nonvolatile storage are on separate controllers with separate power supplies. In case one power supply fails, you have a backup on the other system, as taught on page 30, section 2.7.3 of the IBM reference. The backup storage system of Shibata et al. more closely

resembles the duplex cache writing of the present invention, but fails to mention power management. It would have been obvious to use the system of Beardsley et al. so if a failure in power occurs at any one of the power sources, you will still have a backup of the data in the other memories, and won't lose data, as taught by Beardsley et al. in col. 6, lines 3-8.

29. With respect to claim 13, IBM, Shibata et al., and Kinjo et al. teach all other limitations of the parent claim as discussed supra, but fail to teach the power system of claim 13. Beardsley et al. disclose:

A storage system according to claim 10, further comprising:

a first power unit connected to the first cache memory, in fig. 4, numerals 58 and 82;

a second power unit connected to the second cache memory, the second power unit being independent from the first power unit, in fig. 4, numerals 60 and 96,

Beardsley et al. teach a third cache memory being connected to the second power unit, the third power unit being independent from the first power unit, in fig. 4, numerals 61 and 92, but fail to teach a FIFO buffer. Shibata et al. teach that the third memory can be a FIFO buffer, in col. 4, lines 39-44. Beardsley et al. fail to teach a battery backup. IBM teaches a battery backup on page 89. The NVS, which is the duplex cache memory corresponding to the FIFO in the system of Beardsley, has a battery for feeding power to the NVS in case the power unit fails.

30. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Kinjo et al. and Beardsley et al. before him at the time

the invention was made to modify the backup storage system of IBM, Shibata et al. and Kinjo et al. to include a system of power management as taught by Beardsley et al. With respect to the IBM reference, the first cache memory and the backup nonvolatile storage are on separate controllers with separate power supplies. In case one power supply fails, you have a backup on the other system, as taught on page 30, section 2.7.3 of the IBM reference. The backup storage system of Shibata et al. more closely resembles the duplex cache writing of the present invention, but fails to mention power management. It would have been obvious to use the system of Beardsley et al. so if a failure in power occurs at any one of the power sources, you will still have a backup of the data in the other memories, and won't lose data, as taught by Beardsley et al. in col. 6, lines 3-8. Furthermore, it would have been obvious to one of ordinary skill to also include a battery backup for the FIFO buffer, also to prevent the loss of data in the event of a power failure.

31. Claims 8 and 14 are rejected under 35 U.S.C. 103(a) as being unpatentable over IBM, Shibata et al., Kinjo et al. and Beardsley et al. as applied to claims 1-7, 10-13 and 15 above, and further in view of Yanai et al., US Patent Publication Number 2003/0005355.

32. With respect to claim 8, IBM, Shibata et al., Kinjo et al. and Beardsley et al. teach all other limitations of the parent claims as discussed supra, but fail to include a data indicator. Yanai et al. disclose a storage system wherein the second memory has a unit for indicating presence of absence of data stored in the memory. Referring to the

abstract, when data is written to the first controller, the first write pending indicator is set. Upon completion of the writing of data to the first controller, the first write pending indicator is reset and the second write pending indicator is set. Upon completion of writing data to the second controller, the second write pending indicator is reset.

33. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Beardsley et al., Kinjo et al. and Yanai et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al., Kinjo et al. and Beardsley et al. to include a write pending indicator as taught by the backup storage system of Yanai et al., so data integrity may be maintained, in the event of an error, as taught by Yanai et al. in paragraph 43.

34. With respect to claim 14, IBM, Shibata et al., Kinjo et al. and Beardsley et al. teach all other limitations of the parent claims as discussed supra, but fail to include a data indicator. Yanai et al. disclose a storage system comprising a data remaining indicator for indicating whether or not all the duplicate of the data sent from the host system has been transferred from the memory of the first controller to the memory of the second controller. Referring to the abstract, the second write pending indicator is set when data is being written to the second controller. Upon completion of the writing of data to the second controller, the write pending indicator is reset.

35. It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Beardsley et al., Kinjo et al. and Yanai et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al., Kinjo et al., and Beardsley et al. to include a write pending indicator as

taught by the backup storage system of Yanai et al., so data integrity may be maintained, in the event of an error, as taught by Yanai et al. in paragraph 43.

36. Claim 9 is rejected under 35 U.S.C. 103(a) as being unpatentable over Shibata et al., IBM, Kinjo et al., in view of Matsumoto et al., US Patent 5,720,028.

37. Matsumoto et al. teach a storage system comprising:

a host interface unit connected to a host system, in fig. 2, numeral 210, which is connected to the host system of fig. 1.

a switching unit connected to the host interface unit, in fig. 2 number 220,

a first and a second controller connected to the switching unit, in fig. 2, numerals 200 and 400; and

a storage device connected to the first and the second controller, in fig. 2, numeral 500,

Matsumoto et al. fail to teach the same way to write the memory as described in the present invention. Shibata et al. disclose a first and second controller:

wherein the first controller has a first memory, said first memory being a cache memory, and a second memory, in fig. 1, cache 3 and FIFO 4,

wherein the second controller has a third memory, in fig. 1, numeral 2',

wherein, in the case where the first controller receives data from the host system, the first controller stores the data in the first and the second memory, and then the first controller transfers the data stored in the second memory to the third memory, in col. 4, lines 31-38.

Shibata et al. fail to teach sending a response to the host system, but IBM teaches sending a response to the host system upon receiving data. See page 89 of the IBM publication and page 12 of Applicant's Petition to Make Special.

Shibata fails to teach that the third memory is a cache memory that receives a copy of the data received by said first memory. The third memory of Shibata et al. is a main memory unit. Kinjo et al. teaches that the memory on the second controller can be a cache memory (i.e. the third memory), in col. 6, lines 6-11. This teaches the limitation:

wherein, the second controller has a third memory, said third memory being a cache memory that receives a copy of the data received by said first memory.

38. It would be obvious to one of ordinary skill, having the teachings of IBM and Shibata et al. before him at the time the invention was made, to combine the backup storage system of IBM with the backup storage system of Shibata et al. to eliminate the performance problems due to cross connectional delays, as taught by Shibata et al., and to include response messages as taught by IBM et al. so the host knows when the data has been written to the duplex memories.

Furthermore, it would be obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata, and Kinjo et al. before him at the time the invention was made, to modify the backup storage system of IBM and Shibata with the backup storage system of Kinjo et al. in order to allow high-speed data copy between controllers, as taught by Kinjo et al. in col. 8, lines 61-65. The prior art system of Shibata possibly suffers performance penalties due to the fact that the third memory is a main memory,

so the system of Kinjo et al. resolves this deficiency by using a high-speed cache memory as a third memory.

Furthermore, It would have been obvious to one of ordinary skill in the art, having the teachings of IBM, Shibata et al., Matsumoto et al. and Kinjo et al. before him at the time the invention was made to modify the backup storage system of IBM, Shibata et al., and Kinjo et al. to include the host interface and switching unit of the backup storage system of Matsumoto et al. in order for a specific controller to be disconnected from the larger system, to perform maintenance and management, as taught by Matsumoto in col. 1, line 58 through col. 2 line 4.

Response to Arguments

39. Applicant's arguments, see 14-18, filed 12/14/05, with respect to the rejection(s) of claim(s) 1, 9, 10 and 15 under 35 U.S.C. 103(a) have been fully considered and are persuasive. Therefore, the rejection has been withdrawn. However, upon further consideration, a new ground(s) of rejection is made in view of Kinjo et al. Applicant points out that the present invention writes to a duplexed cache memory on a second controller. This was not the case in the original claims 1 and 9 of the present invention. It was merely claimed that data was written to a memory on the second controller, which the Shibata reference does in fact teach, as it writes to main memory 2' of fig. 1. Applicant amended the claims to more precisely describe that the memory on the second controller is a cache memory to take advantage of the performance benefits of using a fast cache memory as opposed to a main storage device. Kinjo et al. describes

a system where data is written to a cache memory in a first controller. Then a duplicate of the data is transferred to a cache memory in a second controller. Kinjo et al. gives the same motivation for doing this as Applicant, namely to take advantage of the fast cache memory in the second controller, thereby allowing a faster data copy between controllers. Therefore the combination of IBM, Shibata et al. and Kinjo et al. teach all limitations of the amended claim 1 and the original claims 10 and 15. The combination of IBM, Shibata et al., Kinjo et al. and Matsumoto et al. teach all of the limitations of the amended claim 9.

Conclusion

3. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

40. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Ryan Dare whose telephone number is (571) 272-4069. The examiner can normally be reached on Mon-Fri 9:30-6.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matt Kim can be reached on (571) 272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Ryan Dare
January 19, 2006

MATTHEW D. ANDERSON
PRIMARY EXAMINER